

UNITED STATES PATENT APPLICATION

MOBILE STATION DYNAMIC POWER SAVING CONTROL

INVENTORS

**Jiewen Liu
Chih C. Tsien**

Prepared by Dana B. LeMoine
(952) 473-8800

LeMoine Patent Services, PLLC
c/o PortfolioIP
P.O. Box 52050
Minneapolis, MN 55402
ATTORNEY DOCKET 80107.160US1
Client Reference P19061

MOBILE STATION DYNAMIC POWER SAVING CONTROL

Field

5 The present invention relates generally to computer networks, and more specifically to wireless networks.

Background

Wireless networks typically include mobile stations and access points.
10 Mobile stations may run off battery power, and conserving battery power may allow a mobile station to run longer.

Brief Description of the Drawings

Figure 1 shows a diagram of a wireless network;
15 Figures 2 and 3 show flowcharts in accordance with various embodiments of the present invention;
Figures 4 and 5 show desired sleep intervals and wake-up times; and
Figure 6 shows a system diagram in accordance with various embodiments of the present invention.

20

Description of Embodiments

In the following detailed description, reference is made to the accompanying drawings that show, by way of illustration, specific embodiments in which the invention may be practiced. These embodiments are described in sufficient detail to
25 enable those skilled in the art to practice the invention. It is to be understood that the various embodiments of the invention, although different, are not necessarily mutually exclusive. For example, a particular feature, structure, or characteristic described herein in connection with one embodiment may be implemented within other embodiments without departing from the spirit and scope of the invention. In
30 addition, it is to be understood that the location or arrangement of individual elements within each disclosed embodiment may be modified without departing

from the spirit and scope of the invention. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims, appropriately interpreted, along with the full range of equivalents to which the claims are entitled. In the drawings, like numerals
5 refer to the same or similar functionality throughout the several views.

Figure 1 shows a diagram of a wireless network. Wireless network 100 includes access point (AP) 102 and mobile stations (STA) 110, 120, and 130. In some embodiments, wireless network 100 is a wireless local area network (WLAN). For example, one or more of mobile stations 110, 120, and 130, or access point 102
10 may operate in compliance with a wireless network standard such as ANSI/IEEE Std. 802.11, 1999 Edition, although this is not a limitation of the present invention. As used herein, the term “802.11” refers to any past, present, or future IEEE 802.11 standard, including, but not limited to, the 1999 edition.

Mobile stations 110, 120, and 130 may be any type of mobile station capable
15 of communicating in network 100. For example, the mobile stations may be computers, personal digital assistants, wireless-capable cellular phones, or the like. As explained below, in some embodiments, mobile stations 110, 120, and 130 operate in compliance with an 802.11 standard, and are also capable of saving power by entering one of many different possible power savings modes.

20 Access point 102 communicates with mobile station 110 (also referred to as “STA1”) using signal 112. Access point 102 communicates with mobile station 120 (also referred to as “STA2”) using signal 122, and access point 102 communicates with mobile station 130 (also referred to as “STA3”) using signal 132. Mobile station 110 includes network interface card (NIC) 114. In some
25 embodiments, mobile station 110 may be a computer such as a notebook computer or a desktop computer that includes NIC 114. Mobile stations 120 and 130 are shown without NICs. In some embodiments, mobile stations 120 and 130 may be wireless devices that have built-in wireless capability. For example, mobile station 120 may be a notebook computer having a chipset with built-in wireless
30 connectivity capability.

Wireless network 100 is shown as an infrastructure network in which the mobile stations communicate with an access point. In some embodiments of the present invention, mobile stations communicate with other mobile stations directly in an ad hoc network.

5 Mobile stations may save power by “sleeping.” When a mobile station is asleep, all or a portion of the mobile station enters a reduced-power state. Multiple power savings modes may exist. For example, a mobile station may sleep for varying amounts of time to obtain different power saving levels. The remainder of this description details various power saving embodiments of 802.11 compliant
10 mobile stations, but the various embodiments are not limited to 802.11 compliant systems. The discussion is provided within an 802.11 framework to provide examples within a recognizable context, but the various embodiments of the present invention are not meant to be limited to 802.11 systems.

 In some embodiments, an access point may periodically send beacons to
15 mobile stations. For example, access point 102 may periodically transmit beacon frames to be received by any mobile station within range. Access points may also buffer uni-cast data packets for a sleeping station and indicate that in the Traffic Indication Map (TIM) element of beacon. A mobile station in a reduced power state may be required to wake up periodically and listen to the beacons. If the mobile
20 station receives an indication of buffered packets, it may, for example, wake up and poll the packets from the access point. Further, in 802.11 embodiments, broadcast and/or multicast packets that an access point has buffered will be sent out when the Delivery Traffic Indication Message (DTIM) Count in the TIM element becomes zero. The DTIM count is expressed as a number of beacon intervals, and the DTIM
25 count decrements by one in each beacon. A mobile station is also required to wake up for broadcast or multicast packets when the DTIM count becomes zero. The term “broadcast time” as used herein refers to a time at which a mobile station should be awake to receive broadcast or multi-cast packets. In some embodiments, mobile stations include a timer to cause interrupts at Target Beacon Transmission
30 Times (TBTT) based on beacon interval.

Mobile stations may include an Auto-Power-Saving mode (APS). In some embodiments, the APS mode may be enabled or disabled through user configurations. When the APS is disabled, a wireless device is fully powered during operation. While the APS is enabled, the device may enter or exit the sleep mode based on the evaluation of traffic activities. In some embodiments, the traffic activity is a percentage of the time consumed for data traffic over a desired time period. For example, some embodiments include different levels with level five being the deepest sleep mode of power saving modes corresponding to different traffic activities. If transmit/receive (TX/RX) activities in the current time interval are within a range, for example 35%, the device may select power level 2 and enter sleep mode for four beacon intervals. The wake up time for receiving a beacon and evaluation will be at the end of the sleep interval. Table 1, below, shows one example set of relationships among traffic activities, power savings mode, desired sleep interval, and desired wake-up time. In some embodiments, a mobile station may stay awake for a certain time after TX/RX to avoid decision flipping between a sleep mode and normal mode.

Traffic Activities	Power Mode	Desired Sleep Interval	Desired Wake-up Time (for receiving beacon and evaluation)
>60%	Power level 0	Zero	Every beacon interval
40% - 60%	Power level 1	2 beacon intervals	Every 2 beacon intervals
30% - 40%	Power level 2	4 beacon intervals	Every 4 beacon intervals
20% - 30%	Power level 3	6 beacon intervals	Every 6 beacon intervals
10% - 20%	Power level 4	8 beacon intervals	Every 8 beacon intervals
< 10%	Power level 5	10 beacon intervals	Every 10 beacon intervals

Table 1.

The desired sleep interval, expressed in numbers of beacon intervals in Table 1, refers to a desired amount of time that a mobile station may sleep to conserve power until waking up to receive the next beacon, where longer desired sleep intervals correspond to more power conservation. The desired wake-up time
5 describes when a mobile station would wake up based only on the desired sleep interval. As described above, however, mobile stations may need to receive broadcast or multicast packets at broadcast times. Some embodiments determine the actual wake-up time based on the desired sleep interval and the broadcast time. Various embodiments that are described below use both the desired sleep interval
10 and the broadcast time to determine the actual wake-up time.

The example power savings levels shown in Table 1 are shown as discrete levels. In some embodiments, a continuous function is evaluated to determine a next power savings level or a desired sleep interval. The manner in which a power savings level or desired sleep interval is determined as a function of traffic activity
15 is not a limitation of the present invention.

Figure 2 shows a flowchart in accordance with various embodiments of the present invention. In some embodiments, method 200 describes the operation of a mobile station when not asleep, also referred to as “normal mode.” In some
20 embodiments, method 200, or portions thereof, is performed by a mobile station, a network interface card, a processor, or an electronic system, embodiments of which are shown in the various figures. Method 200 is not limited by the particular type of apparatus, software element, or system performing the method. The various actions in method 200 may be performed in the order presented, or may be performed in a different order. Further, in some embodiments, some actions listed in Figure 2 are
25 omitted from method 200.

Method 200 is shown beginning at block 210 in which the mobile station is operating normally with Auto-Power savings (APS) mode enabled. In some
embodiments, when APS is not enabled, method 200 is not performed, and the mobile station stays powered on. At 215, if a host requests that all or a portion of
30 the mobile station be turned off or put to sleep, control is transferred to 255, which

is described further below. In some embodiments, the host referred to at 215 may be a processor within a mobile station, and the actions of 215 may be performed by a network interface card or a processor coupled to a wireless network interface.

At 220, if the current time interval has expired, method 200 continues to 225
5 to check if a beacon is received. If either the current time interval has not expired, or if a beacon is not received, method 200 stays in normal operation at 210. The time interval referred to at 220 may refer to any wake-up time interval selected based on traffic activities. In some embodiments, the current time interval may be expressed in a number of beacon intervals. For example, the current time interval
10 may be one beacon interval, two beacon intervals, or more.

At 230, traffic activity is evaluated. In some embodiments, the traffic activity is measured as a percentage of the current time interval that is used for data traffic. For example, the volume of transmit/receive (TX/RX) traffic in the current time interval may be measured, and that may be compared against the length of the
15 current time interval. In some embodiments, traffic activity is evaluated using measurements other than total data traffic volume. For example, traffic activity may be expressed in terms of any metric available to the mobile station.

At 235, a next power savings level is selected and a desired sleep interval is selected. In some embodiments, the next power savings level and the desired sleep
20 interval may be selected as a function of the traffic activity evaluated at 230. Referring now back to Table 1, the next power savings level and desired sleep interval may be selected from the table based on traffic activity as listed in the left hand column. For example, if the data traffic exceeds 60% of the current time interval, the next power savings level is selected as power level 0. In some
25 embodiments, the threshold may be higher or lower than 60%. In other embodiments, the next power savings level and the desired sleep interval are selected in a manner different than that represented by Table 1. Further, in some embodiments, the next power savings level may be selected as a function of the traffic activity in the current time interval as well as one or more previous time
30 intervals.

At 240, if the next power savings level selected corresponds to power level 0, method 200 remains in normal operation mode at 210. At 245, if power saving conditions are OK, method 200 proceeds with 250, otherwise method 200 remains in normal operation mode at 210. Powers saving conditions may be, for example,

5 related to criteria such as:

- whether the NIC is associated with an AP;
- whether the transmission path is empty for certain time period;
- whether the receiving path is empty;
- whether the station is in the middle of process, such as messages
- 10 exchanging between stations/AP or between host and NIC,
- scanning, configuration changing, etc.; or
- whether radio parts that will be switched off or switched to sleep
- mode are ready to be switched accordingly.

If any power savings condition fails at 245, then method 200 may remain in normal
15 operation at 210.

At 250, a wake-up time is calculated. The wake-up time referred to at 250 refers to the time at which the mobile station will wake up after being asleep. In some embodiments, the wake-up time corresponds to the end of the desired sleep interval. In various embodiments of the present invention, the next wake-up time is
20 determined based on the desired sleep interval and the broadcast time. Example embodiments to determine the next wake-up time are described more fully below with reference to Figures 4 and 5.

At 255, the apparatus performing method 200 notifies the AP that it is going to sleep. In some embodiments, the apparatus performing method 200 is a mobile
25 station such as mobile station 110 (Figure 1), and access point 102 is notified that mobile station 110 is going to sleep. At 260, the radio in the wireless interface is powered down and the station is prepared for sleep mode, and at 265, the station enters sleep mode. In some embodiments, only the transmitter of the radio is powered down to allow for continued reception, and in other embodiments, the
30 transmitter and receiver are both powered down. When in sleep mode, a mobile

station may have every subsystem put in a reduced power state, or may have less than every subsystem put in a reduced power state. For example, in some embodiments, a network interface card (NIC) within a mobile station performs method 200, and only a portion of the NIC is put into a reduced power state when
5 asleep.

Figure 3 shows a flowchart in accordance with various embodiments of the present invention. In some embodiments, method 300 describes the operation of a mobile station when asleep to save power. In some embodiments, method 300, or portions thereof, is performed by a mobile station, a network interface card, a
10 processor, or an electronic system, embodiments of which are shown in the various figures. Method 300 is not limited by the particular type of apparatus, software element, or system performing the method. The various actions in method 300 may be performed in the order presented, or may be performed in a different order. Further, in some embodiments, some actions listed in Figure 3 are omitted from
15 method 300.

Method 300 is shown beginning at block 310 in which the mobile station is in sleep mode. The mobile station may enter sleep mode as shown in method 200 (Figure 2). If the host interrupts the apparatus performing method 300 at 315, or if the mobile station has missed too many beacons at 320, method 300 continues with
20 block 350 to transition to normal mode. Block 350 is described further below. In some embodiments, the host referred to at 315 may be a processor within a mobile station, and the actions of 315 may be performed by a network interface card or a processor coupled to a wireless network interface.

If not interrupted by the host, and if too many beacons have not been missed,
25 method 300 continues at 325. At 325, method 300 determines if it is time to wake up for a beacon. If not time to wake up, method 300 loops in sleep mode. An apparatus may wake up for a beacon by responding to an interrupt by a TBTT timer. For example, a timer may be set to a wake-up time prior to the apparatus going to sleep, and when the timer expires, the hardware may be interrupted, and a beacon
30 monitor program or task may be performed. If, at 330, a beacon is received, then

the next wake-up time is adjusted at 335 to attempt to obtain a longer sleep time period, and if a beacon is not received, the next wake-up time is set to one beacon interval in an attempt to receive the next beacon.

In some embodiments, if a received beacon includes an indication of buffered traffic, a mobile station or NIC performing method 300 may poll an access point to receive the traffic. For example, between 330 and 335, data traffic may be exchanged between a mobile station and an access point. When the data traffic is complete, the next wake-up time is adjusted at 335, and the mobile station returns to sleep mode at 310.

As described above, method 300 may transition to normal mode if the host desires to transmit, or if too many beacons have been missed. At 350, the next wake-up time is set to the end of the next beacon interval, the radio is powered up at 355, the AP is notified at 360, and normal mode is entered at 365.

In some embodiments, all or portions of methods 200 (Figure 2) and 300 (Figure 3) are implemented in a task that is run each time a beacon is supposed to arrive at a mobile station. For example, a beacon monitor may be a task or a procedure that keeps running during normal operation to track incoming beacons. The task may be notified using a timer interrupt such as a TBTT interrupt for every incoming beacon. An interrupt driven process such as this may be able to provide accurate timing for power saving control. As a control point set in the beacon monitor, the power saving control procedures may be called each time a beacon is supposed to arrive. The control procedure may then perform all related processes, such as traffic activity evaluations, decisions regarding entering or exiting sleep mode, conditions checking prior to entering sleep mode, hardware powering off or on, and the like.

Figures 4 and 5 show sequences that demonstrate possible scenarios when determining a wake-up time based on a desired sleep interval and a broadcast interval (the period between broadcast times). At the end of a desired sleep interval, a mobile station in power saving mode may wake up for receiving a beacon, while at the end of a broadcast interval it may also wake up for broadcast or multicast

packets. Figures 4 and 5 show the calculation to determine the wake up time. For example, Figure 4 shows wake-up times when the desired sleep interval is greater than the broadcast interval. Sequence 400 shows a desired sleep interval of six beacon intervals. The desired sleep interval may be determined based on the traffic evaluation during the current time interval. For example, in embodiments represented by Table 1 above, a device enters power level three if traffic activity is less than 30% of the current time interval, and the desired sleep interval will be selected as six beacon intervals.

Also shown in Figure 4 is a broadcast interval equal to four beacon intervals. As described above, the broadcast interval may correspond to the interval between times at which a device needs to wake up to receive broadcast and/or multicast packets. In 802.11 embodiments, this may be at the DTIM time. If the desired sleep interval is greater than the broadcast interval (as is the case in Figure 4), the next wake-up time may be set to the next broadcast time.

Figure 5 shows a sequence in which the desired sleep interval is less than the broadcast interval. As shown in sequence 500, the desired sleep interval is two beacon intervals, and the broadcast interval is five beacon intervals. In some embodiments, the next wake-up time is set to the end of the desired sleep interval if at least two sleep intervals exist before the next broadcast time. In other words, the next wake-up time is set to the end of the desired sleep interval if the broadcast time is more than two desired sleep intervals since the last wake-up time. For example, at 502, the next broadcast time is five beacon intervals in the future, which is more than twice the desired sleep interval of two, so the next wake-up time is set to the end of a desired sleep interval, which is at 504. At 504, the next broadcast time is three beacon intervals away, which is less than twice the desired sleep interval, so the next wake-up time is set to be equal to the broadcast time at 506. For the example of Figure 5, this pattern continues, and wake-up times are alternately scheduled for two and three beacon intervals until either the desired sleep interval or the broadcast interval is changed.

The calculations to determine the next wake-up time as demonstrated in Figure 5 may also be expressed with the following pseudo-code:

```

    If ((the difference between the next desired wake-up time and next
5    broadcast time) > desired sleep interval) Then
        Set the next wake-up time to the next desired wakeup time
    Else
        Set the next wake-up time to the next broadcast time.

10    For example, at 502, the next desired wake-up time is at 504 (2 beacon
    intervals away) and the next broadcast time is at 506 (5 beacon intervals away). The
    difference between 504 and 506 is three beacon intervals, which is greater than the
    desired sleep interval (two beacon intervals). Therefore, the next wake-up time is
    set to 504 (this is the “if” case in the pseudo-code example above). At 504, the next
15    desired wake-up time is at 508 and the next broadcast time is at 506. The difference
    between 508 and 506 is one beacon interval, which is less than the desired sleep
    interval. Therefore the next wake-up time is set to 506 (this is the “else” case in the
    pseudo-code example above).
```

Figures 4 and 5 represent methods to determine a next wake-up time based
20 on a desired sleep interval and a broadcast interval, where the desired sleep interval
is based on an amount of traffic in a current interval. In some embodiments, the
next wake-up time is determined in a manner different than that represented in
Figures 4 and 5. For example, in some embodiments, a mobile station may need to
be awake for more than one broadcast time, or for various other events. This
25 information may be combined in any manner with a desired sleep interval that is
derived from a traffic volume without departing from the scope of the present
invention.

Figure 6 shows a system diagram in accordance with various embodiments
of the present invention. Electronic system 600 includes antennas 610, radio
30 interface 620, physical layer (PHY) 630, media access control (MAC) mechanism

640, processor 660, and memory 670. In some embodiments, electronic system 600 may be a mobile station, a wireless interface, a NIC, or the like. For example, electronic system 600 may be utilized in network 100 as any of mobile stations 110, 120, or 130, or NIC 114. Also for example, electronic system 600 may be an
5 apparatus capable of performing any of the method embodiments described with reference to the previous figures.

In some embodiments, electronic system 600 may represent a system that includes a wireless interface as well as other circuits. For example, in some embodiments, electronic system 600 may be a computer, such as a personal
10 computer, a workstation, or the like, that includes a wireless interface as a peripheral or as an integrated unit.

In operation, system 600 sends and receives signals using antennas 610, and the signals are processed by the various elements shown in Figure 6. Antennas 610 may include one or more directional antennas or one or more omni-directional
15 antennas. As used herein, the term omni-directional antenna refers to any antenna having a substantially uniform pattern in at least one plane. For example, in some embodiments, antennas 610 may include an omni-directional antenna such as a dipole antenna, or a quarter wave antenna. Also for example, in some embodiments, antennas 610 may include a directional antenna such as a parabolic
20 dish antenna or a Yagi antenna. In some embodiments, antennas 610 form an array capable of supporting spatial division multiple access (SDMA) or multiple-input multiple output (MIMO) communications. In other embodiments, antennas 610 include only one physical antenna.

Radio interface 620 is coupled to antennas 610 to interact with a wireless
25 network. Radio interface 620 may include circuitry to support the transmission and reception of radio frequency (RF) signals. For example, in some embodiments, radio interface 620 includes an RF receiver to receive signals and perform “front end” processing such as low noise amplification (LNA), filtering, frequency conversion or the like. Further, in some embodiments, radio interface 620 includes
30 beamforming circuitry to support SDMA processing. Also for example, in some

embodiments, radio interface 620 includes circuits to support frequency up-conversion, and an RF transmitter. The invention is not limited by the contents or function of radio interface 620.

5 All or a portion of radio interface 620 may include the capability to be put into a reduced power state. For example, a transmitter within radio interface 620 may include circuitry to allow the transmitter to be powered-down or partially powered down. In some embodiments the transmitter may be powered down during periods that electronic system 600 is asleep. Further, in some embodiments, all of radio interface 620 may be powered down when electronic system 600 is asleep.

10 Physical layer (PHY) 630 may be any suitable physical layer implementation. For example, PHY 630 may be a circuit block that implements a physical layer that complies with an IEEE 802.11 standard or other standard. Examples include, but are not limited to, direct sequence spread spectrum (DSSS), frequency hopping spread spectrum (FHSS), and orthogonal frequency division
15 multiplexing (OFDM). Like radio interface 620, all or a portion of PHY 630 may be put into a reduced power state to support various power savings modes.

Media access control (MAC) mechanism 640 may be any suitable media access control layer implementation. For example, MAC 640 may be implemented in software, or hardware or any combination thereof. In some embodiments, a
20 portion of MAC 640 may be implemented in hardware, and a portion may be implemented in software that is executed by processor 660. Further, MAC 640 may include a processor separate from processor 660.

Processor 660 may perform method embodiments of the present invention, such as method 200 (Figure 2) or method 300 (Figure 3). Processor 660 represents
25 any type of processor, including but not limited to, a microprocessor, a digital signal processor, a microcontroller, or the like.

Memory 670 represents an article that includes a machine readable medium. For example, memory 670 represents a random access memory (RAM), dynamic random access memory (DRAM), static random access memory (SRAM), read only
30 memory (ROM), flash memory, or any other type of article that includes a medium

readable by processor 660. Memory 670 may store instructions for performing the execution of the various method embodiments of the present invention.

Although the present invention has been described in conjunction with certain embodiments, it is to be understood that modifications and variations may be
5 resorted to without departing from the spirit and scope of the invention as those skilled in the art readily understand. Such modifications and variations are considered to be within the scope of the invention and the appended claims.